

# An Evaluation of Sparse Inverse Covariance Models for Group Functional Connectivity in Molecular Imaging

D.B.. Keator<sup>a,1,\*</sup>, A. Ihler<sup>b,2</sup>

<sup>a</sup>University of California, Irvine, Irvine Hall rm. 163 - ZOT 3960, Irvine CA. 92697-3435

<sup>b</sup>University of California, Irvine, Donald Bren Hall, Irvine CA. 92697-3435

---

## Abstract

Evaluating the functional relationships between brain regions measured with neuroimaging provides insight into how the brain is sharing information at a macro scale. Many functional connectivity methods have been developed for dynamic imaging modalities such as functional MRI (fMRI), while less work has focused on models for static molecular imaging techniques such as FDG-PET and Tc-99m HMPAO SPECT across groups of individuals. In this work we provide a quantitative assessment of how well three functional connectivity models based on sparse inverse covariance estimation can accurately recover gold standard connectivity patterns across multiple cohorts and data set sizes. We compare the accuracies of learning regularized inverse covariance matrices across cohorts independently with those learned using two different group-based regularization models. By using large cohorts of SPECT scans, we are able to provide a quantitative assessment of the accuracy of the models in recovering the gold standard functional connectivity patterns and connectivity strengths across a range of typical data set sizes in neuroimaging studies.

*Keywords:* Functional connectivity, SPECT, PET, graphical model

---

## 1. Introduction

Neuroimaging modalities are routinely used to evaluate regional differences in brain function and connectivity. Neuroimaging provides researchers with a method to evaluate in-vivo brain function, where various statistics relating activity between spatially distributed regions are used as indicators of functional connectivity [1]. Many studies have shown changes in functional connectivity, as assessed with neuroimaging, across genders [2, 3] and in psychiatric illnesses such as schizophrenia [4, 5] and Alzheimer’s Disease [6, 7] among others. The majority of studies and methods of evaluating functional connectivity have focused on dynamic imaging modalities such as functional magnetic resonance imaging (fMRI) and 15O-H<sub>2</sub>O Positron Emission Tomography (PET). Fewer studies have evaluated functional connectivity in static molecular imaging modalities such as <sup>18</sup>F-FDG PET and Tc-99m HMPAO Single Photon Emission Computed Tomography (SPECT) across groups of subjects. Huang et al. [6] compares functional connectivity across cohorts using

---

\*Corresponding author

<sup>1</sup>Email Address: dbkeator@uci.edu (D.B. Keator)

<sup>2</sup>Email Address: aihler@uci.edu (A. Ihler)

static PET; however, their results give little information on how well the models recover the true connectivity profiles of the populations tested. Careful validation of connectivity models is necessary, prior to their use and interpretation in real data [8].

A relatively simple method of interrogating functional connectivity in neuroimaging is through inter-regional and/or voxel-wise correlations [9, 10]. This method typically consists of cross-correlating the mean functional activity for each region of interest (ROI) and applying a threshold to remove values close to zero. Unfortunately, it has been shown that the sample correlation matrix is often a poor estimator of the population correlation matrix and can result in incorrect conclusions based on estimation error [11] [8]. Further, correlations between two ROIs does not imply a direct functional connection and thus impedes our ability to interpret our results with respect to known brain circuits. A more interpretable statistic, that has been shown to estimate the true network [8], is the partial correlation: the statistical relationship between pairs of brain regions after conditioning on the others. If the brain regions are jointly Gaussian, then we can find regions with strong partial correlation by examining the non-zero entries in the scaled inverse covariance matrix of the Gaussian distribution, fit to the observed data. Unfortunately, in typical neuroimaging studies, we do not have enough data to accurately compute the inverse of the covariance matrix among brain regions; therefore, some form of regularization is needed. To estimate the inverse covariance matrix and thus the partial correlations consistent with our observed data, we can use regularized maximum likelihood methods [12, 13]. Although we have methods available, we lack data on how well these classes of models can recover the true connectivity in group-based molecular imaging and how sample sizes effect accuracy of the results.

In this study we compare three models for learning functional connectivity: one where the connectivity profiles are learned independently for each cohort and two that share information between cohorts. In settings with low sample sizes, we hypothesize that models which share connectivity information across cohorts will perform better in recovering correct connections. Our contribution is a quantitative assessment of how well the functional connectivity models can accurately recover a gold standard connectivity pattern across a range of typical data set sizes in neuroimaging studies based on a very large sample of SPECT scans. To our knowledge, this is the first study evaluating the accuracy of recovered connectivity profiles by sample size in group-based static molecular imaging. The rest of the manuscript is organized as follows. In the section 2, we describe three models for learning sparse inverse covariance matrices which have been applied to similar problems. Next, we describe the gold standard connectivity data used to quantitatively evaluate the inverse covariance models. We conclude section 2 with a summary of the model. In section 3, we present a detailed comparison of the sparse inverse covariance models with respect to a gold standard connectivity profile, as a function of sample size using two different methods of setting the model regularization coefficients. We conclude the manuscript with a discussion of the results and plans for future work.

## 2. Methods

In this section we describe three sparse inverse covariance models for learning group-based functional connectivity patterns between pairs of regions, one of which learns the patterns independently for each cohort and two models which share connectivity information across cohorts. We anticipate the models that share connectivity information across cohorts will improve our prediction accuracy in settings with small sample sizes. We conclude the methods section with a description of the gold standard dataset we use to quantitatively evaluate the performance of each connectivity model.

### 2.1. *Sparse Inverse Covariance Estimation*

Ultimately we are interested in learning about how the brain is functionally connected at the group level across different experimental conditions and/or cohorts. In learning about the functional connectivity in our data, we are presented with a few challenges. First, the brain is composed of approximately 100 billion neurons, with each neuron being connected with up to 10,000 other neurons, potentially forming 1 trillion synaptic connections. Even at the scale of neuroimaging, we are likely to find some association between any pair of brain regions. We need a principled approach that focuses our attention on the most relevant associations in a way that can be related to known brain circuits. Next, we are constrained by small sample sizes. Often, particularly in rare diseases, it is difficult to find large groups of subjects, making it more difficult to learn about whole-brain functional connectivity. Fortunately, there are models available that can help us learn about functional connectivity in light of these problems, yet there is little data on how accurately the models perform when applied to molecular brain imaging data.

Sparse inverse covariance estimation models provide a structured approach of estimating pairwise connections between variables in settings with small numbers of samples. Entries in the inverse covariance matrix (also called the precision matrix) of a Gaussian model, correspond to the pairwise statistical relationships between variables conditioned on all the other variables in the model. Values of zero in the inverse covariance matrix (i.e., zero partial correlation) imply conditional independence of the associated variables [27]. Alternatively, strong partial correlations are suggestive of pairwise interactions, helping us to interpret results with respect to known brain circuits. In contrast, bivariate correlations provide pairwise associations but do not account for all the other brain regions and their potential effect on the correlation; although, using both partial and bivariate correlations together when interpreting results is well advised. Learning inverse covariance matrices in settings with small data set sizes and a large number of variables is difficult and some form of regularization is needed to both improve prediction accuracy and aid in interpretation, focusing on a smaller subset that exhibit the strongest effects [12]. In section 2.1.1 we describe Graphical Lasso (GL), a popular technique for learning sparse inverse covariance matrices. GL uses regularization to improve prediction accuracy and

interpretability in settings with small sample sizes relative to the number of variables [12]; however, in many experiments the groups being compared are similar, and we would expect some of the functional connectivity relationships are shared across cohorts. We would like to use this observation to help improve our prediction accuracy. In section 2.1.2 we describe the Fused Graphical Lasso (FGL) and Group Graphical Lasso (GGL) models proposed by Danaher et al. [28] based on work from Tibshirani et al. [29] for jointly estimating sparse precision matrices across cohorts. Our contribution is a quantitative analysis of how well these models perform in recovering a gold standard connectivity profile as a function of sample size in static molecular imaging.

### 2.1.1. Graphical Lasso

The least absolute shrinkage and selection operator (lasso) proposed by Tibshirani [12] is a method for penalized regression that shrinks some parameter estimates and sets others to zero. Shrinking parameter estimates controls over fitting and results in models that have better predictive accuracy. Setting some parameters to zero provides variable selection, focusing on the more relevant features of the data, resulting in more interpretable modeling results. Banerjee et al. [13] applied the lasso penalty to learning sparse undirected graphical models in a multivariate Gaussian setting and developed a block-wise interior point algorithm which they noted is equivalent to iteratively solving lasso problems. Friedman et al. [30] pursued this observation and developed the Graphical Lasso (GL) algorithm. The learning problem is to maximize the penalized log-likelihood over all positive definite matrices  $\Phi$  given by:

$$\hat{\Sigma}^{-1} = \arg \max_{\Phi \succ 0} \log \det \Phi - \text{tr}.(S\Phi) - \lambda \|\Phi\|_1 \quad (1)$$

where  $\hat{\Sigma}^{-1}$  is the estimate of the sparse precision matrix,  $S$  is the empirical covariance matrix of the data computed using regions of interest (see Appendix A), where entry  $S_{i,j}$  is the covariance between regions  $i$  and  $j$  across all subjects in the group, and  $\|\Phi\|_1$  is the  $L_1$  norm, the sum of the absolute values of the elements of the current estimate of  $\Phi$ . The non-negative coefficient  $\lambda$  controls the relative importance of the sparsity-inducing  $L_1$  regularization term. To apply the GL algorithm to our problem, we compute the covariance matrix between all pairs of regions for each cohort separately and directly apply the GL algorithm to learn a precision matrix for each cohort independently. The GL algorithm cycles through the variables, fitting a modified lasso regression to each variable in turn. The algorithm is fast, enabling problems with thousands of parameters to be solved efficiently. In estimating functional connectivity across the entire brain, we can easily have thousands of parameters and small sample sizes, making this algorithm attractive.

### 2.1.2. Joint Graphical Lasso

In many neuroimaging experiments we are searching for subtle changes in functional relationships between a subset of brain areas, i.e., we are not expecting the relationships between all brain regions to be different

across cohorts. In situations where we expect parameters to be shared across groups and have a small data set size relative to the number of parameters being estimated, it has been shown by Danaher et al. [28], using simulated data with a known amount of parameter sharing, that jointly estimated sparse precision matrices are closer to the true distributions than estimating them separately with the GL algorithm. The learning problem is similar to equation (1) except we replace the  $L_1$  norm penalty with a generic penalty function  $P(\{\Phi\})$  defined across the set of precision matrices  $\{\Phi\}$ :

$$\left[ \left( \hat{\Sigma}^{(1)} \right)^{-1}, \dots, \left( \hat{\Sigma}^{(G)} \right)^{-1} \right] = \arg \max_{\Phi \succ 0} \left( \sum_{g=1}^G n_g \left[ \log \det \Phi^{(g)} - \text{tr}(S^{(g)} \Phi^{(g)}) \right] - P(\{\Phi\}) \right) \quad (2)$$

where  $\left[ \left( \hat{\Sigma}^{(1)} \right)^{-1}, \dots, \left( \hat{\Sigma}^{(G)} \right)^{-1} \right]$  are the estimates of the sparse precision matrices for all groups,  $n_g$  is the number of observations in group  $g$ , and the first term inside the sum is the contribution from each group to the log-likelihood. Danaher et al. [28] propose two convex penalty functions, the Fused Graphical Lasso (FGL) and Group Graphical Lasso (GGL) penalties. FGL encourages groups to share both network structure and parameter values; the FGL penalty is given by:

$$P(\{\Phi\}) = \lambda_1 \sum_{g=1}^G \sum_{i \neq j} \left| \Phi_{i,j}^{(g)} \right| + \lambda_2 \sum_{g < g'} \sum_{i,j} \left| \Phi_{i,j}^{(g)} - \Phi_{i,j}^{(g')} \right| \quad (3)$$

where  $\lambda_1$  and  $\lambda_2$  are non-negative coefficients controlling the amount of sparsity and similarity in the precision matrices across groups respectively. The first term in equation (3) is similar to the  $L_1$  penalty in graphical lasso, except that the sum is over the off-diagonal elements. Increasing this penalty results in sparser networks. The second term is the absolute difference in corresponding parameter estimates across the  $G$  precision matrices that we are estimating, one for each cohort. Increasing the  $\lambda_2$  penalty will result in similar connectivity networks and strengths between pairs of variables across the groups. Making  $\lambda_2$  extremely large will result in identical networks across the groups.

In contrast, the GGL penalty encourages a shared pattern of sparsity, without requiring similarity between the parameter estimates across the groups. The GGL objective is:

$$P(\{\Phi\}) = \lambda_1 \sum_{g=1}^G \sum_{i \neq j} \left| \Phi_{i,j}^{(g)} \right| + \lambda_2 \sum_{i \neq j} \left( \sum_{g=1}^G \Phi_{i,j}^{(g)^2} \right)^{\frac{1}{2}} \quad (4)$$

where  $\lambda_1$  and  $\lambda_2$  are non-negative coefficients controlling the amount of sparsity and shared sparsity respectively. The first term in equation (3) is the same  $L_1$  penalty as FGL; as before, making this penalty large results in sparser networks. The second term encourages a similar pattern of sparsity in the precision matrices across the groups. Interestingly, both penalty terms will contribute to sparsity resulting in a weaker form of network similarity than FGL, where patterns of sparsity (e.g., functional connectedness) are similar across the groups, but the strength of connections in the network are not penalized for being different.

To estimate the networks using joint graphical lasso penalties, we use the alternating directions method

of multipliers (ADMM) [32]. The algorithm is guaranteed to converge to the global optimum. For details on applying ADMM to this problem see Danaher et al. [28]. To apply the joint graphical lasso algorithms to our problem, we follow the same procedure as for graphical lasso except the precision matrices are learned jointly across cohorts. Compared to graphical lasso, the joint lasso penalties have two coefficients to set,  $\lambda_1$  and  $\lambda_2$ . In section 3 we will evaluate different ways of setting the penalties and whether the joint models improve accuracy in recovering the gold standard connectivity patterns and strengths across a range of data set sizes.

## 2.2. Gold Standard Data

Ultimately we are interested in learning about differences in how the brain is sharing information across groups of subjects, evaluated with static molecular imaging techniques, because it may provide important information on overall brain function. In practical settings, the amount of available data is small relative to the number of parameters needed to estimate whole-brain connectivity, which could result in both spurious connections and/or poor estimates of true connections due to dependencies in the parameters [11]. In section 2.1 we described three models for functional connectivity that use regularization to, hopefully, improve predictive accuracy and reduce variance in our parameter estimates, at the expense of some bias, while also making the results more interpretable, focusing our attention on the strongest functional connections [12]. In order to evaluate the performance of the inverse covariance models in finding true functional connections in molecular imaging data, we require a large data set with many samples.

The Amen Clinics Inc. (<http://www.amenclinics.com/>) has been collecting technetium-99m hexamethylpropyleneamine oxide (Tc99m HMPAO) SPECT scans for many years on a variety of disorders such as attention deficit hyperactivity disorder, depression, anxiety, and behavioral problems, among others. The Amen Clinics Inc. provided us with a SPECT data set to use as a gold standard, consisting of 11,906 males (avg. age  $29.2 \pm 17.6$ ) and 7,550 females (avg. age  $35.6 \pm 18.4$ ) [33]. This is the largest SPECT data set that we know of. On average, 55% of the subjects have mood disorders, 7% bipolar, 25% depression, 47% ADHD, 45% anxiety, 10% substance abuse, and 32% brain trauma. Although the data set consists of a variety of disorders and comorbidities, for this purpose we need two or more large cohorts that have both shared and unshared functional connections, to evaluate how well the three inverse covariance models recover these patterns as a function of sample size.

### 2.2.1. Regions of Interest

Prior to performing an analysis of functional connectivity, we must define the nodes of our network. In our gold standard dataset, structural scans were unavailable and thus a popular choice is to use a standard

atlas to define regions of interest (ROIs) such as the AAL atlas from Tzourio-Mazoyer et al. [14] and average the functional voxels contained in each anatomical boundary. Unfortunately, this approach essentially ignores the functional data in defining the ROIs and is prone to decreased sensitivity by averaging signals from a small number of activated voxels with noise [15]. Here we take a pragmatic approach, given the data we have available for evaluating the inverse covariance models, and use Gaussian finite mixture modeling to find functional clusters constrained by anatomically-defined boundaries for use as nodes in the network. Our motivation is to find a set of regions that are consistent with the observed functional data across all subjects, while also allowing the nodes to be interpreted with respect to an anatomical reference. Details about applying Gaussian finite mixture modeling to define regions of interest can be found in Appendix A. Briefly, for each anatomical region defined in the AAL atlas, we use the observed counts for each voxel contained in the anatomical region, across all subjects, and perform Gaussian finite mixture modeling to learn clusters of functional activity. We then use the Bayesian information criterion (BIC) score to select the number of clusters (i.e.,  $K$ ) for each region that best fits our observed data [25]. Once the number of clusters have been selected for each anatomical region, we compute values for each cluster by averaging over the number of detections at each voxel, weighted by their cluster membership probabilities, for each subject separately. We use these cluster averages, computed for each subject, in the sparse inverse covariance models.

### 2.2.2. Scan Preparation

Subjects in the Amen data set were injected with an age/weight appropriate dose of Tc99m HMPAO and were at rest during uptake. All subjects were scanned on a high-resolution Picker Prism 3000 triple-headed gamma camera with fan beam collimators. The original reconstructed image matrices were 128x128x29 voxels with sizes of 2.16mm x 2.16mm x 6.48mm and values representing counts. The images were spatially normalized to the MNI atlas using SPM8 software [34], resulting in image matrices of 79 x 95 x 68 voxels in x, y, and z dimensions respectively with isotropic 2mm voxel sizes. The Automated Anatomical Labeling (AAL) atlas was used to define the brain regions based on the anatomical parcellations available in the atlas because no structural imaging data was available [14].

After normalization, the clustering model (see Appendix A) for  $K \in \{1, \dots, 25\}$  was run across the entire data set, for each region in the AAL atlas, except the cerebellum. The cerebellum was not included due to missing data in the lower slices of some scans. The clustering model, after selecting the best model for each anatomical region using the BIC scores, resulted in 180 clusters across the brain. The cluster averages for each subject, stratified by cohorts, were computed and mean centered. Although the sample is large relative to the number of variables, there is still measurement noise in the data. To remove some of this noise, we run the graphical lasso algorithm with light regularization ( $\lambda = 0.1$ ), resulting in precision matrices with 2119 pairwise connected regions in males (partial correlation range -0.66-0.17), 2130 pairwise connections

in females (partial correlation range -0.64-0.17), and 1750 shared connections ( $\approx 82\%$ ) across both cohorts. These precision matrices will be used in section 3 as a gold standard, to quantitatively evaluate the performance of the inverse covariance estimation models in recovering the patterns as a function of the amount of available training data.

### 2.3. Model Summary

The clustering and sparse inverse covariance algorithms are summarized graphically in Figure 1. The entire set of imaging data is collapsed across cohorts prior to clustering. For each region, `kmeans++` is run for each setting of  $K$  and used to initialize a Gaussian mixture model [23]. The resulting clusters are evaluated using the BIC scores and the solution with the minimum BIC score for each region are selected and used as regions of interest (see Appendix A). The region of interest averages are then calculated for each subject and each cluster across all regions. For each cohort, we then compute the empirical covariance matrix among all pairs of clusters and learn precision matrices using the graphical lasso, fused, and group graphical lasso algorithms. This process results in inverse covariance matrices with varying amounts of non-zero entries which are depicted in a graph structure by adding an edge between pairs of variables. Figure 2 shows an example of clustering and the graphical lasso model run with eight anatomical regions (i.e., yellow region outlines) where voxels have been colored by their cluster membership within each anatomical region. The edges in the graph indicate non-zero entries in the inverse covariance matrix between pairs of clusters and are scaled by the magnitude of the partial correlations, where blue edges indicate negative partial correlations and red areas positive partial correlations. Pairs of clusters that are not connected in the figure have zero entries in the inverse covariance matrix, reflecting their conditional independence. Note, in the experiments discussed in section 3, these procedures were run in three-dimensions using the whole brain.

## 3. Results

In this section we provide a quantitative evaluation of the sparse inverse covariance models in terms of correct connections and connection strengths, across different sample sizes compared to the gold standard precision matrices described in section 2.2. The intent of our experiments are to evaluate the performance of the inverse covariance models at sample sizes typically used in functional imaging studies. Each of the models described in section 2.1 requires a setting for the amount of regularization, which affects the sparsity of the results and the overall network structure; therefore, we evaluate two methods of setting the regularization weights. In section 3.0.1 we use the gold standard precision matrices directly to determine optimal settings for the  $\lambda_2$  regularization coefficients in the FGL and GGL models and compare receiver operating characteristic (ROC) curves of true-positive (TP) and false-positive (FP) connections across all models by varying  $\lambda$  in GL and  $\lambda_1$  in FGL and GGL. In practice, one does not typically have access to a gold standard. In section 3.1 we use cross-validation to determine the regularization weights and compare the resulting precision matrices with those from the gold standard.



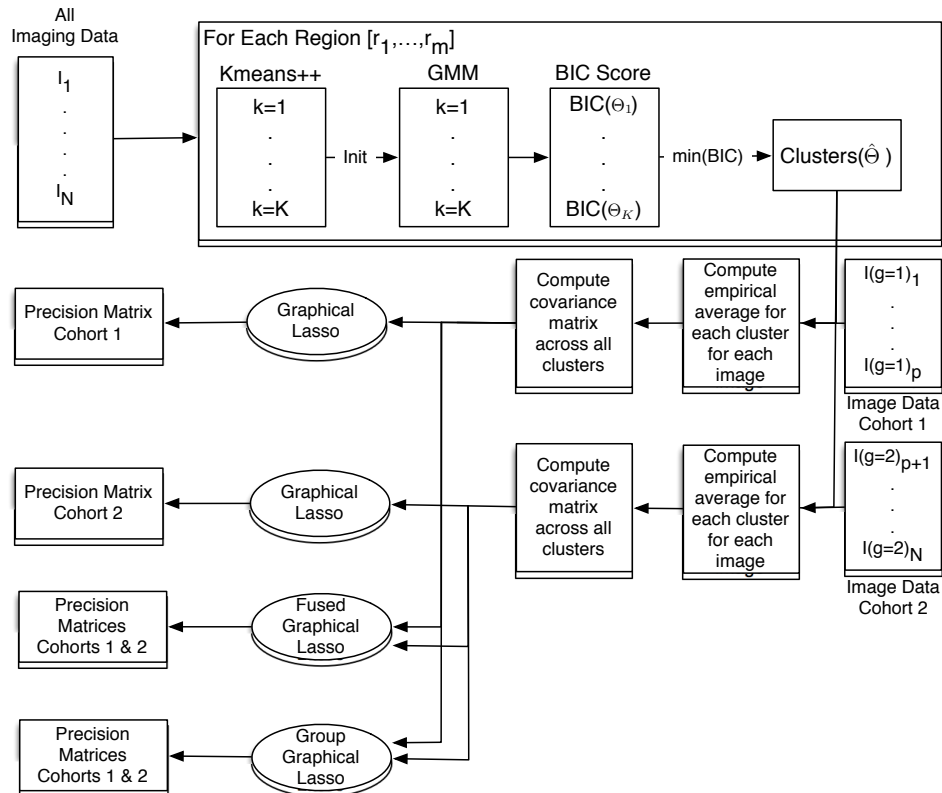


Figure 1: Summary of the methods for learning functional clusters and precision matrices using graphical lasso, fused graphical lasso, and group graphical lasso algorithms in group-based static molecular imaging data. Note, in this example there are only two cohorts but the process is the same for additional groups.




Figure 2: Example of clustering and inverse covariance modeling in a two-dimensional SPECT scan slice. Yellow outlines correspond to anatomical regions from an atlas, voxels are colored by their cluster membership within each anatomical region, and lines connecting clusters correspond to the non-zero parameters in the inverse covariance matrix between pairs of clusters and are scaled by the magnitude of the partial correlations, where blue edges indicate negative partial correlations and red areas positive partial correlations.

### 3.0.1. Using Gold Standard for Parameter Settings

In this experiment we evaluate the sparse inverse covariance models by sample size, using the gold standard precision matrices to determine settings for the regularization parameters. To compare GL, FGL, and GGL models, we follow a procedure similar to Danaher et al. [28]. We fix the regularization coefficients of the group penalties (i.e.,  $\lambda_2$ ) and vary the  $\lambda_1$  penalty in the FGL and GGL models and the  $\lambda$  penalty in the GL model, creating ROC curves of model accuracy in predicting correct functional connections. Further, we compare the strength of connections relative to the gold standard precision matrices. To determine fixed settings for the group penalties, we use the following procedure. We first draw two random samples ( $n=500$ ) from the gold standard data set, stratified by group. We then compute the cluster averages using the clusters found with the gold standard data set described in section 2.2. Using these data, we perform a grid search over the  $\lambda_1$  and  $\lambda_2$  parameter space, for the FGL and GGL models separately, comparing the learned precision matrices ( $\hat{\Phi}$ ) to those of the gold standard ( $\Phi^*$ ) in terms of their  $L_1$  difference  $\left(\sum_{g=1}^G \left|\hat{\Phi}_g - \Phi_g^*\right|\right)$ . We select the settings that result in the minimum  $L_1$  difference from the gold standard. For the FGL model, the minimum was found with settings  $\lambda_1 = 0.01$  and  $\lambda_2 = 0.1$ . For the GGL model, the minimum was found with settings  $\lambda_1 = 0.001$  and  $\lambda_2 = 0.01$ . We will use these fixed settings for  $\lambda_2$  in our comparison of the group regularized models (i.e., FGL and GGL) with GL.

Next, we draw 10 subsets of images, with replacement, stratified by groups, with equal sizes  $N \in \{10, 25, 50, 100, 250, 500, 750, 1000, 1500, 2000, 2500\}$  from the gold standard data set, after removing the subjects used in the parameter searching routine described above. Using the clusters found for the gold standard data set described in section 2.2, we compute the cluster averages and mean center the data. Next, we learn precision matrices using the GL model across a wide range of  $\lambda$  regularization penalties. Similarly, using the fixed  $\lambda_2$  settings, we learn precision matrices using the FGL and GGL models across the same range of  $\lambda_1$  regularization settings as with GL.

The results showing the average number (across both groups and 10 subsets) of true-positive edges (TP) correctly identified (i.e., non-zero entries in the test precision matrices match those from the gold standard), compared to the average number of false-positive edges (FP) incorrectly identified (i.e., non-zero entries in the test precision matrices unmatched with the gold standard) by sample sizes are shown in Figures 3-5. The colored regions around each marker indicate the unit standard deviation region about the means. The large markers identify points on the curves with the lowest sum of squared errors (SSE) between the test precision matrices and the gold standard for each model, indicating the point on the curves where connection strengths in the test precision matrices best match the gold standard. Interestingly, these occur in regions with a high proportion of true-positive edges relative to false positive edges. Visually, the FGL and GGL model curves generally dominate the GL curves for most sample sizes. The biggest difference between the joint models and graphical lasso are in the moderate sample sizes  $N \in \{25, 50, 100\}$ , where we see the largest benefit from

parameter sharing across the groups. Both joint models, in general, perform equally well in terms of the presence or absence of edges. The average areas under the TP/FP edge curves (AUC) across the subsets, by sample sizes, are shown in Table 1 (columns 2-4). The T-test and Bonferroni corrected p-values, comparing all combinations of the models by sample sizes, are shown in columns 5-7. The AUCs are statistically lower in the GL model (i.e., negative t-scores) than both the FGL and GGL models at data set sizes less than or equal to 1000, with the moderate sample sizes showing the largest decreases. At data set sizes above 2000, the GL model outperforms the joint models, which is related to having enough data to accurately estimate the precision matrices without the added benefit of parameter sharing in the joint models, and the fact that we used the GL algorithm to remove noise when creating the gold standard precision matrices.

To evaluate how well the entries in the precision matrices corresponding to the strength of functional connections match those of the gold standard data set, we calculate the SSE between the test precision matrices and the gold standard by model and sample size. The average minimum SSEs (across both groups and 10 subsets) are shown in Table 2 (columns 2-4). The T-test and Bonferroni corrected p-values comparing all combinations of the models by sample sizes are shown in Table 2 (columns 5-7). The results are similar to the TP/FP edge results, where the joint models perform better than the GL model at moderate sample sizes, with the largest difference at  $N = 100$ . Between the FGL and GGL joint models, the FGL model performs better, yielding lower SSEs compared to the gold standard. Unlike the TP/FP edge results, the FGL model achieves a lower SSE than the GL model, even for the larger data sets (i.e  $N \in \{1500, 2000, 2500\}$ ), whereas GGL does not. To evaluate the effect of the clustering model on these results, we repeated our experiments by averaging all the voxels within each region of interest instead of using the clustering model. The results from the three sparse inverse covariance models are shown in Appendix Tables B.5 and B.6 and are consistent with those presented in this section.

### 3.1. Using Cross-Validation for Parameter Settings

In section 3.0.1 we used the precision matrices from the gold standard data set directly in determining fixed group regularization penalties (i.e.,  $\lambda_2$ ) for the joint models. In practice, one does not typically have access to a gold standard. In this experiment we consider how well these models perform in a practical setting, determining the regularization coefficient settings using a typical 10-fold cross validation design. For this experiment we require 10 subsets of data, stratified by group, for regularization coefficient selection and testing with the gold standard. We first separate the gold standard data set into two independent sets, stratified by group ( $N^{(g1)} = 5,953$ ,  $N^{(g2)} = 3,775$  in each set). From each set and group, we draw 10 subsets of images, with replacement, of equal sizes  $N \in \{10, 25, 50, 100, 250, 500\}$ . This procedure results in 10 subsets of data for regularization coefficient selection orthogonal to the 10 subsets for precision matrix evaluation. Using the clusters learned from the gold standard dataset, we compute the average values for each cluster and each subject in the corresponding regularization coefficient selection subset and compute the corresponding covariance matrices between each of the clusters. We now have 20 covariance matrices, one for each of the 10

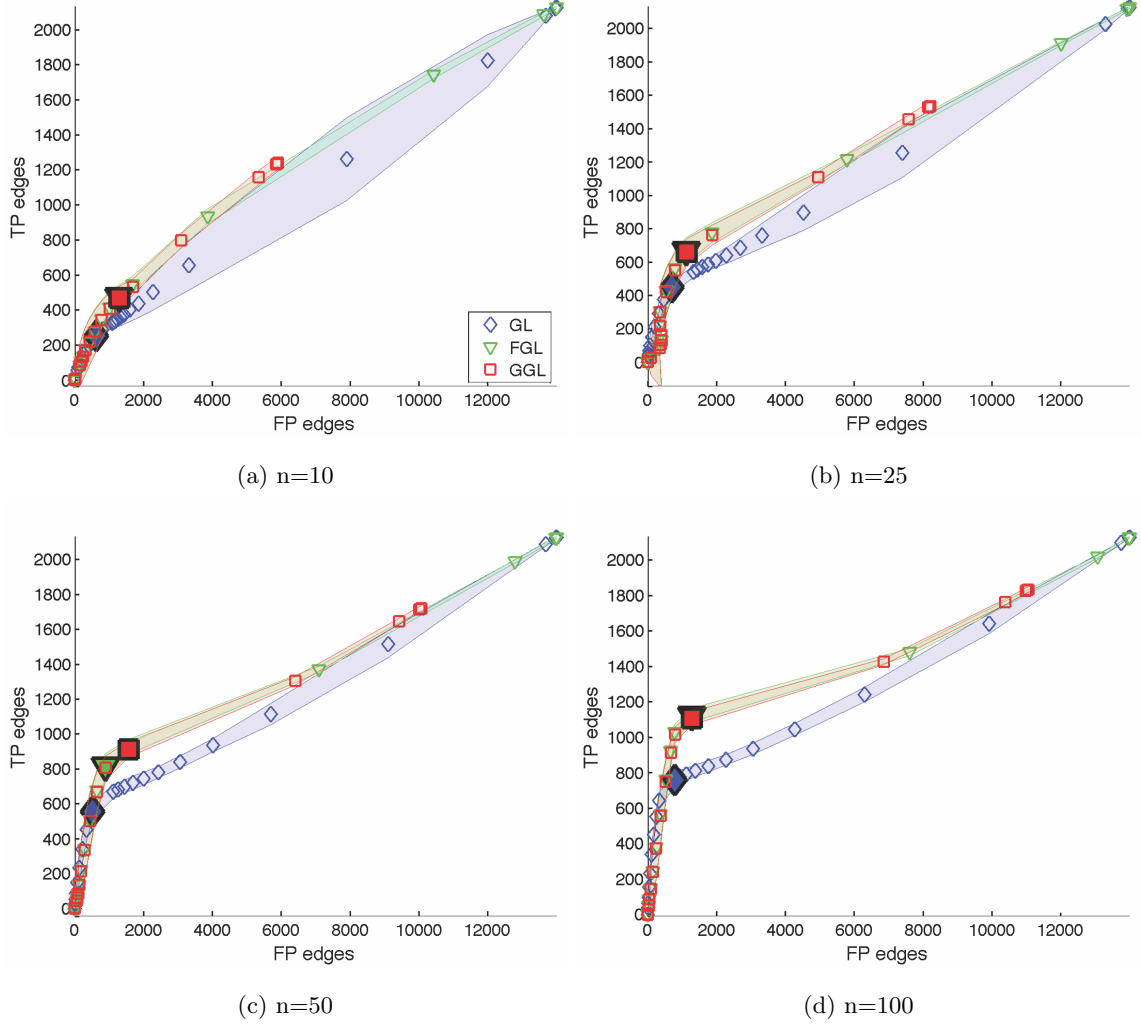





Figure 3: True-positive (TP) vs. false-positive (FP) edges between the test precision matrices and the gold standard for sample sizes  $n \in \{10, 25, 50, 100\}$  using GL, FGL, and GGL inverse covariance models. Small markers indicate the means across the 10 test subsets and both cohorts. Colored regions show the unit standard deviation about the means. Large markers identify points with lowest average SSE compared to gold standard.




(a)  $n=250$



(b)  $n=500$



(c)  $n=750$



(d)  $n=1000$

Figure 4: True-positive (TP) vs. false-positive (FP) edges between the test precision matrices and the gold standard for sample sizes  $n \in \{250, 500, 750, 1000\}$  using GL, FGL, and GGL inverse covariance models. Small markers indicate the means across the 10 test subsets and both cohorts. Colored regions show the unit standard deviation about the means. Large markers identify points with lowest average SSE compared to gold standard.

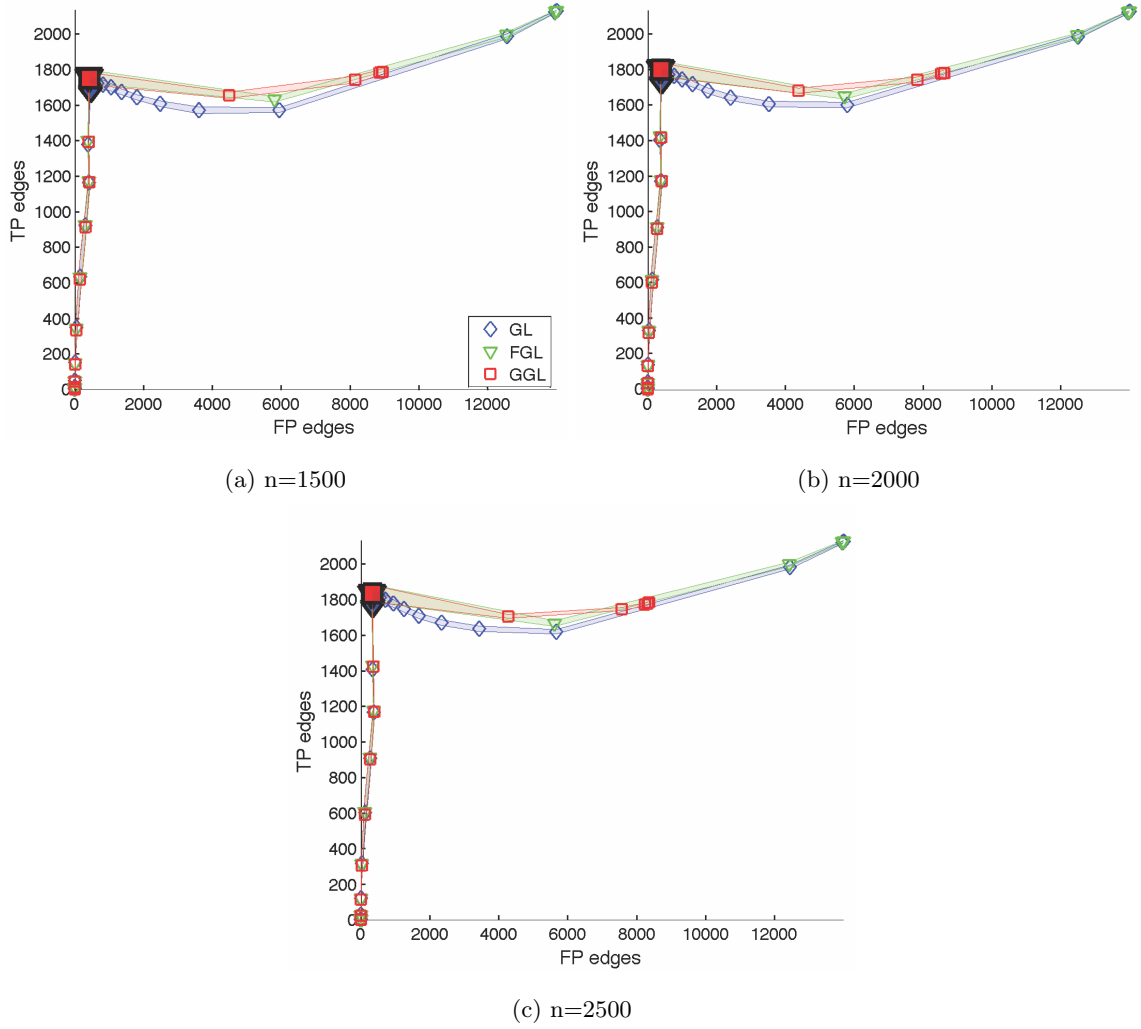


Figure 5: True-positive (TP) vs. false-positive (FP) edges between the test precision matrices and the gold standard for sample sizes  $n \in \{1500, 2000, 2500\}$  using GL, FGL, and GGL inverse covariance models. Small markers indicate the means across the 10 test subsets and both cohorts. Colored regions show the unit standard deviation about the means. Large markers identify points with lowest average SSE compared to gold standard.

Table 1: TP/FP Edge Area Under Curve (AUC) T-test Comparison by Model and Sample Size using Gold Standard for Parameter Settings

	GL	FGL	GGL	GL-FGL	GL-GGL	FGL-GGL
N	$\overline{AUC} \pm std$	$\overline{AUC} \pm std$	$\overline{AUC} \pm std$	T(P Bonferroni)	T(P Bonferroni)	T(P Bonferroni)
2500	0.820 $\pm$ 0.003	0.790 $\pm$ 0.008	0.810 $\pm$ 0.008	10.74( <b>1.15E-07</b> )	3.93( <b>3.79E-02</b> )	-5.51( <b>1.22E-03</b> )
2000	0.812 $\pm$ 0.004	0.797 $\pm$ 0.015	0.806 $\pm$ 0.013	2.88(3.87E-01)	1.31(2.08E-01)	-1.33(1.00e00)
1500	0.799 $\pm$ 0.004	0.792 $\pm$ 0.018	0.806 $\pm$ 0.013	1.19(1.00e00)	-1.44(1.00e00)	-1.87(1.00e00)
1000	0.784 $\pm$ 0.004	0.799 $\pm$ 0.005	0.811 $\pm$ 0.003	-8.14( <b>7.52E-06</b> )	-18.18( <b>1.93E-11</b> )	-6.70( <b>1.09E-04</b> )
750	0.769 $\pm$ 0.004	0.789 $\pm$ 0.007	0.797 $\pm$ 0.005	-8.26( <b>6.07E-06</b> )	-14.16( <b>1.31E-09</b> )	-2.79(4.68E-01)
500	0.683 $\pm$ 0.006	0.776 $\pm$ 0.006	0.779 $\pm$ 0.004	-36.99( <b>7.65E-17</b> )	-45.47( <b>1.93E-18</b> )	-1.27(1.00e00)
250	0.660 $\pm$ 0.003	0.746 $\pm$ 0.007	0.745 $\pm$ 0.007	-35.21( <b>1.83E-16</b> )	-35.30( <b>1.75E-16</b> )	0.13(1.00e00)
100	0.631 $\pm$ 0.006	0.696 $\pm$ 0.006	0.695 $\pm$ 0.006	-24.07( <b>1.50E-13</b> )	-23.26( <b>2.73E-13</b> )	0.21(1.00e00)
50	0.604 $\pm$ 0.004	0.652 $\pm$ 0.005	0.653 $\pm$ 0.005	-22.95( <b>3.44E-13</b> )	-24.43( <b>1.16E-13</b> )	-0.45(1.00e00)
25	0.572 $\pm$ 0.004	0.621 $\pm$ 0.013	0.626 $\pm$ 0.012	-11.28( <b>5.28E-08</b> )	-13.19( <b>4.26E-09</b> )	-0.85(1.00e00)
10	0.537 $\pm$ 0.007	0.601 $\pm$ 0.013	0.604 $\pm$ 0.013	-13.24( <b>3.98E-09</b> )	-14.30( <b>1.12E-09</b> )	-0.58(1.00e00)

regularization coefficient selection subsets, for each cohort, using the clusters learned from the gold standard data. These covariance matrices will be used to determine settings for the regularization coefficients. To set the regularization coefficients, we perform grid searches over the regularization parameter spaces (i.e.,  $\lambda_1$ ,  $\lambda_2$  for FGL and GGL, and  $\lambda$  for GL) for each inverse covariance model, selecting the regularization coefficient settings that minimizes the BIC score from each regularization coefficient selection subset, where the number of parameters are the number of non-zero entries in the precision matrices. After performing this procedure, we have settings for the regularization coefficients for each of the three inverse covariance models, for each of the 10 subsets of data. Finally, we use the gold standard clusters and the regularization coefficients to compute the precision matrices for each of the three inverse covariance models across the 10 precision matrix evaluation subsets and both cohorts.

The results from this experiment are shown in Figure 6 and Tables 3 and 4. Figure 6 and Table 3 show the average number of TP and FP connections, by model, across the 10-folds compared to the gold standard data set. The lower segment of the stacked bars (i.e., lighter color) shows the number of FP connections whereas the upper segment (i.e., darker color) shows the number of TP connections. For  $N \in \{100, 250, 500\}$  the FGL and GGL models find more TP connections than GL. For  $N \in \{10, 25, 50\}$  GL finds more TP connections than the joint models yet all of the models are performing poorly. The SSE between the average precision matrices across the 10 test subsets and the gold standard are shown in Table 4. The results are similar to those of the TP/FP edges for  $N \in \{100, 250, 500\}$ , where FGL and GGL models achieve sta-



Table 2: Sum of Squared Errors (SSE) with Gold Standard by Model and Sample Size

	GL	FGL	GGL	GL-FGL	GL-GGL	FGL-GGL
N	$\overline{SSE} \pm std$	$\overline{SSE} \pm std$	$\overline{SSE} \pm std$	T(P Bonferroni)	T(P Bonferroni)	T(P Bonferroni)
2500	1.35 $\pm$ 0.27	1.34 $\pm$ 0.18	1.74 $\pm$ 0.16	0.11(1.00E+00)	-3.93( <b>3.90E-02</b> )	-5.26( <b>3.90E-03</b> )
2000	1.66 $\pm$ 0.23	1.57 $\pm$ 0.20	2.08 $\pm$ 0.30	0.96(1.00E+00)	-3.56(8.58E-02)	-4.54( <b>1.17E-02</b> )
1500	2.16 $\pm$ 0.25	1.90 $\pm$ 0.21	2.33 $\pm$ 0.22	2.52(8.35E-01)	-1.66(1.00E+00)	-4.55( <b>7.80E-03</b> )
1000	3.29 $\pm$ 0.25	2.70 $\pm$ 0.41	3.09 $\pm$ 0.30	3.85( <b>4.68E-02</b> )	1.59(1.00E+00)	-2.44(9.95E-01)
750	4.32 $\pm$ 0.37	3.47 $\pm$ 0.43	3.92 $\pm$ 0.53	4.77( <b>7.80E-03</b> )	1.94(1.00E+00)	-2.11(1.00E+00)
500	6.03 $\pm$ 0.68	4.78 $\pm$ 0.67	5.15 $\pm$ 0.44	4.14( <b>2.34E-02</b> )	3.42(1.17E-01)	-1.49(1.00E+00)
250	11.41 $\pm$ 1.34	8.92 $\pm$ 1.14	8.68 $\pm$ 0.76	4.46( <b>1.17E-02</b> )	5.61( <b>3.90E-03</b> )	0.56(1.00E+00)
100	67.51 $\pm$ 18.79	22.58 $\pm$ 1.69	21.20 $\pm$ 1.40	7.53( <b>3.90E-03</b> )	7.77( <b>3.90E-03</b> )	1.98(1.00E+00)
50	64.34 $\pm$ 7.46	60.42 $\pm$ 6.83	57.70 $\pm$ 4.71	1.23(1.00E+00)	2.38(1.00E+00)	1.04(1.00E+00)
25	62.98 $\pm$ 5.59	58.46 $\pm$ 5.74	62.70 $\pm$ 5.84	1.78(1.00E+00)	0.11(1.00E+00)	-1.64(1.00E+00)
10	100.52 $\pm$ 10.29	83.90 $\pm$ 4.48	84.13 $\pm$ 4.49	4.68( <b>7.80E-03</b> )	4.62( <b>7.80E-03</b> )	-0.11(1.00E+00)

tistically lower SSEs than the GL model. In contrast, at the lower  $N$ 's, there is no significant difference between the models. Interestingly, the average SSE measurements for the GL model in the larger groups (i.e.,  $N \in \{100, 250, 500\}$ ) are much higher than in the experiment described in section 3.0.1; whereas, in the FGL and GGL models the results are more consistent. In looking at the regularization parameter settings across the splits of the data for this experiment, we find the GL model is selecting higher regularization parameters (range 0.2 – 0.7 depending on split and number of samples) as compared to the settings yielding the minimum average SSE when using the gold standard to set them directly (range 0.1 – 0.3). Conversely, in the FGL and GGL models, we find the  $\lambda_1$  regularization parameter settings to be similar to those found yielding the minimum average SSE when using the gold standard; whereas, the  $\lambda_2$ , group regularization setting is generally a bit higher. Because the group models use all the data, across both cohorts, in setting the regularization parameters, the settings are more stable across splits and result in better estimates of the connection strengths in sample sizes  $N \in \{100, 250, 500\}$  as demonstrated in Table 4, yet in the smallest sample sizes this is not the case.

#### 4. Discussion

In section 3 we described two experiments to evaluate the inverse covariance models described in section 2. In section 3.0.1 we used the gold standard precision matrices to set the group regularization parameters in the joint models that yielded the smallest difference from the gold standard. The results suggest that the joint models perform better in detecting true-positive functional connections and more accurately learn the connection strengths across most of the sample sizes. At  $N \in \{1500, 2000, 2500\}$  the joint models find




Figure 6: TP/FP edges by sample size and model for cross-validation experiment. Lower, light colored segments show the average number of FP edges whereas, upper, dark colored segments show the average number of TP edges.

Table 3: TP/FP Edges by Model and Sample Size Using Cross-Validation

	GL		FGL		GGL	
N	$\overline{TP} \pm \text{std}$	$\overline{FP} \pm \text{std}$	$\overline{TP} \pm \text{std}$	$\overline{FP} \pm \text{std}$	$\overline{TP} \pm \text{std}$	$\overline{FP} \pm \text{std}$
500	1090.6 $\pm$ 193.1	464.1 $\pm$ 153.6	1540.1 $\pm$ 48.1	641.0 $\pm$ 127.2	1625.2 $\pm$ 24.0	1236.4 $\pm$ 170.6
250	937.9 $\pm$ 159.4	517.6 $\pm$ 146.5	1400.2 $\pm$ 35.3	752.8 $\pm$ 42.9	1274.4 $\pm$ 35.9	903.4 $\pm$ 45.1
100	613.2 $\pm$ 215.5	585.8 $\pm$ 257.2	1195.8 $\pm$ 30.6	1058.2 $\pm$ 62.2	1062.2 $\pm$ 43.6	1223.8 $\pm$ 141.7
50	483.5 $\pm$ 346.0	322.1 $\pm$ 288.2	497.2 $\pm$ 409.0	377.6 $\pm$ 372.2	242.3 $\pm$ 285.2	215.3 $\pm$ 313.6
25	239.5 $\pm$ 193.4	281.6 $\pm$ 267.9	91.0 $\pm$ 111.0	61.8 $\pm$ 114.2	69.0 $\pm$ 107.9	50.9 $\pm$ 98.1
10	106.8 $\pm$ 207.7	292.4 $\pm$ 642.1	66.7 $\pm$ 148.1	117.6 $\pm$ 316.1	67.9 $\pm$ 146.1	123.0 $\pm$ 317.0

Table 4: Sum of Squared Errors (SSE) with Gold Standard by Model and Sample Size Using Cross-Validation

	GL	FGL	GGL	GL-FGL	GL-GGL	FGL-GGL
N	$\overline{SSE} \pm \text{std}$	$\overline{SSE} \pm \text{std}$	$\overline{SSE} \pm \text{std}$	T(P Bonferroni)	T(P Bonferroni)	T(P Bonferroni)
500	156.66 $\pm$ 70.11	4.88 $\pm$ 1.04	8.54 $\pm$ 3.32	6.85( <b>3.77E-05</b> )	6.67( <b>5.26E-05</b> )	-3.33(6.76E-02)
250	190.47 $\pm$ 41.17	7.16 $\pm$ 0.96	11.46 $\pm$ 1.56	14.08 ( <b>6.69E-10</b> )	13.74( <b>9.98E-10</b> )	-7.42( <b>1.26E-05</b> )
100	207.17 $\pm$ 46.93	14.28 $\pm$ 1.64	31.82 $\pm$ 16.17	13.00( <b>2.51E-09</b> )	11.20( <b>2.85E-08</b> )	-3.41(5.59E-02)
50	266.49 $\pm$ 75.68	231.15 $\pm$ 128.5	332.05 $\pm$ 91.61	0.75(1.00E+00)	-1.74(1.00E+00)	-2.02(1.00E+00)
25	283.22 $\pm$ 68.10	366.5 $\pm$ 55.47	372.8 $\pm$ 49.69	-3.00(1.39E-01)	-3.36(6.27E-02)	-0.27(1.00E+00)
10	530.03 $\pm$ 167.41	342.01 $\pm$ 41.8	350.91 $\pm$ 45.33	3.45(5.19E-02)	3.27(7.73E-02)	-0.46(1.00E+00)

$\sim 2 - 3\%$  more true positive edges than the GL algorithm, at  $N \in \{250, 500, 750, 1000\}$  the joint models find  $\sim 3 - 7\%$  more true positive edges than GL, and at  $N \leq 100$  the joint models find  $\sim 31 - 46\%$  more true positive edges than the GL algorithm. The largest gain for the joint models (i.e., FGL and GGL) over the independent model (i.e., GL) in our comparisons were at  $N = 100$ . In section 3.1 we used a cross-validation design and orthogonal splits of the gold standard data to evaluate the performance of these models in a more practical setting, where a gold standard is not available. The results of this experiment similarly show the joint models performing better than the GL model for data set sizes of 100 or more, with the largest difference at  $N = 100$ . At  $N \in \{100, 250, 500\}$  the joint models find  $\sim 26 - 59\%$  more true positive edges than the GL algorithm. At  $N = 50$  the FGL and GL algorithms detect only  $\sim 23\%$  of the total true positive edges; whereas, the GGL algorithm is much lower at  $\sim 11\%$  of the true positive edges. Even worse, at  $N \leq 25$  the true positive rates for all models are  $\sim 3 - 11\%$  which is far too low for practical use. Further, we observe that the GL model SSE measures are much higher across the range of data set sizes when using cross validation to set the regularization weight than when using the gold standard directly (section 3.0.1), caused by over regularization of the GL model in the cross-validation experiment. The results at sample sizes less than 50 are generally quite poor and none of the models perform well using cross-validation and

BIC scores to select settings for the regularization parameters. Furthermore, this result is consistent when using cross-validation and the Akaike Information Criterion (AIC) score, a score that penalizes less for the number of model parameters than BIC, to select the regularization parameters. Visually, we see the GL model with more true-positive edges at  $N \in \{10, 25\}$  than the joint models; although, all models appear to be over regularizing. From the results of experiment 1 (see section 3.0.1), using the gold standard to set the regularization parameters, we know the joint models can perform better at the smallest sample sizes, recovering  $\sim 21 - 22\%$  of the true positive edges at  $N = 10$  versus  $\sim 11\%$  for the GL model, yet it is evident that care should be taken when setting the regularization parameters in the cross-validated design with such small samples. It may be prudent to use a larger hold-out sample to set the regularization parameters prior to applying them to very small datasets. Alternatively, one could compute the inverse covariance models across a range of regularization parameters and interpret the results based on edge selection frequencies or by controlling for false discovery rates as in Liu et al. [35].

Our results generally agree with Danaher et al. [28] who showed an improvement in the accuracy of the joint models over graphical lasso in groups with shared patterns of edges in simulated data. In our experiments, using molecular imaging data, both joint models perform well with the FGL model moderately outperforming the GGL model in some experiments. In practice, one should take into consideration the expected similarity of the groups being compared and the amount of available data, to determine whether the added complexity of selecting two regularization coefficients in the FGL and GGL models is warranted and to choose among the joint models.

## 5. Conclusions

In this manuscript we have compared models for group-based functional connectivity using static molecular imaging data and given a quantitative evaluation of the models in recovering a gold standard connectivity profile, as a function of sample size. In the largest samples, all models performed well; although, the joint models generally perform better. In smaller samples, the joint models are more stable and achieve better results when using a large dataset to determine settings for the regularization coefficients. Caution should be used when applying these models to very small data sets. Our experiments suggest the true positive rates will be low and the connection strengths will be inaccurate when using cross-validation and BIC scores to set the regularization parameters. Given these caveats, our results show there is value in using sparse inverse covariance estimation for measuring functional connectivity in group-based molecular imaging. It would be interesting to extend these experiments to within-subject experimental designs. In recent work, Qiu et al. [36] has developed a joint penalty that appears to perform better than GGL in settings where there are dependencies between networks. Using experiments similar to those presented in this manuscript, an evaluation by sample size of within-subject designs would provide complimentary information on using sparse inverse covariance estimation for functional connectivity modeling in molecular imaging.

## **6. Acknowledgements**

For this work we wish to thank Dr. Daniel Amen, M.D. and the Amen Clinics Inc. for supplying the gold standard dataset. This work was supported in part by the National Center for Research Resources and the National Center for Advancing Translational Sciences, National Institutes of Health, through Grant UL1 TR000153. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

## 7. References

### References

1. Friston K. Functional and effective connectivity in neuroimaging: a synthesis. *Human brain mapping* 1994;2(1-2):56–78.
2. Ingalhalikar M, Smith A, Parker D, Satterthwaite TD, Elliott MA, Ruparel K, Hakonarson H, Gur RE, Gur RC, Verma R. Sex differences in the structural connectome of the human brain. *Proceedings of the National Academy of Sciences* 2014;111(2):823–8.
3. Biswal BB, Mennes M, Zuo XN, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, et al. Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences* 2010;107(10):4734–9.
4. McGuire P, Frith C. Disordered functional connectivity in schizophrenia. *Psychological medicine* 1996;26(04):663–7.
5. Micheloyannis S, Pachou E, Stam CJ, Breakspear M, Bitsios P, Vourkas M, Erimaki S, Zervakis M. Small-world networks and disturbed functional connectivity in schizophrenia. *Schizophrenia research* 2006;87(1):60–6.
6. Huang S, Li J, Sun L, Liu J, Wu T, Chen K, Fleisher A, Reiman E, Ye J. Learning brain connectivity of alzheimer’s disease from neuroimaging data. In: *NIPS*; vol. 22. 2009:808–16.
7. Wang K, Liang M, Wang L, Tian L, Zhang X, Li K, Jiang T. Altered functional connectivity in early alzheimer’s disease: A resting-state fmri study. *Human brain mapping* 2007;28(10):967–78.
8. Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, Ramsey JD, Woolrich MW. Network modelling methods for fmri. *Neuroimage* 2011;54(2):875–91.
9. Rogers BP, Morgan VL, Newton AT, Gore JC. Assessing functional connectivity in the human brain by fmri. *Magnetic resonance imaging* 2007;25(10):1347–57.
10. Clark C, Kessler R, Buchsbaum M, Margolin R, Holcomb H. Correlational methods for determining regional coupling of cerebral glucose metabolism: A pilot study. *Biological Psychiatry* 1984;.
11. Varoquaux G, Craddock RC. Learning and comparing functional connectomes across subjects. *NeuroImage* 2013;80:405–15.
12. Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)* 1996;:267–88.
13. Banerjee O, El Ghaoui L, d’Aspremont A. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research* 2008;9:485–516.

14. Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *Neuroimage* 2002;15(1):273–89.
15. Poldrack RA. Region of interest analysis for fmri. *Social cognitive and affective neuroscience* 2007;2(1):67–70.
16. Hirai T, Jones E. A new parcellation of the human thalamus on the basis of histochemical staining. *Brain Research Reviews* 1989;14(1):1–34.
17. Saper C, Chelimsky T. A cytoarchitectonic and histochemical study of nucleus basalis and associated cell groups in the normal human brain. *Neuroscience* 1984;13(4):1023–37.
18. Mai JK, Assheuer J, Paxinos G. Atlas of the human brain. Academic Press San Diego; 1997.
19. McLachlan GJ, Basford KE. Mixture models. inference and applications to clustering. *Statistics: Textbooks and Monographs, New York: Dekker, 1988* 1988;1.
20. Toncray JE, Krieg WJ. The nuclei of the human thalamus: a comparative approach. *Journal of Comparative Neurology* 1946;85(3):421–59.
21. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society Series B (Methodological)* 1977;:1–38.
22. McLachlan G, Krishnan T. The EM algorithm and extensions; vol. 382. John Wiley & Sons; 2007.
23. Arthur D, Vassilvitskii S. k-means++: The advantages of careful seeding. In: *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics; 2007:1027–35.
24. Smyth P. Model selection for probabilistic clustering using cross-validated likelihood. *Statistics and Computing* 2000;10(1):63–72.
25. Schwarz G, et al. Estimating the dimension of a model. *The annals of statistics* 1978;6(2):461–4.
26. Haughton DM, et al. On the choice of a model to fit data from an exponential family. *The Annals of Statistics* 1988;16(1):342–55.
27. Marrelec G, Krainik A, Duffau H, Péligrini-Issac M, Lehericy S, Doyon J, Benali H. Partial correlation for functional brain interactivity investigation in functional mri. *Neuroimage* 2006;32(1):228–37.
28. Danaher P, Wang P, Witten DM. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2014;76(2):373–97.

29. Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2005;67(1):91–108.
30. Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 2008;9(3):432–41.
31. Hsieh CJ, Sustik MA, Dhillon I, Ravikumar P, Poldrack R. Big & quic: Sparse inverse covariance estimation for a million variables. In: *Advances in Neural Information Processing Systems*. 2013:3165–73.
32. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 2011;3(1):1–122.
33. Amen D, Trujillo M, Keator D, Taylor D, Willeumier K. Gender differences in regional cerebral blood flow in a healthy and psychiatric cohort of 46,034 spect scans. *The Open Neuroimaging Journal* Under Review;.
34. Penny WD, Friston KJ, Ashburner JT, Kiebel SJ, Nichols TE. Statistical Parametric Mapping: The Analysis of Functional Brain Images: The Analysis of Functional Brain Images. Academic Press; 2011.
35. Liu W, et al. Gaussian graphical model estimation with false discovery rate control. *The Annals of Statistics* 2013;41(6):2948–78.
36. Qiu H, Han F, Liu H, Caffo B. Joint estimation of multiple graphical models from high dimensional time series. *arXiv preprint arXiv:13110219* 2013;.



## Appendix A. Functional Clusters

Here we describe a finite mixture model for finding functional clusters within anatomically-defined regions of interest that will subsequently be used as nodes in our network for functional connectivity modeling. The model has been designed with respect to static SPECT and PET modalities, consisting of independent subjects across multiple cohorts, with no available anatomical scans. Our motivation is to find a set of regions that are consistent with the observed functional data across all cohorts, while also allowing the nodes to be interpreted with respect to an anatomical reference. From neuroanatomy, we know the brain is fundamentally composed of spatially varying clusters of cells which segregate during development into specialized units. These specialized units consist of compact clusters of neurons, forming nuclei, which contribute to the overall function of the region. Neuroanatomic studies of brain cytoarchitectonics provides evidence for intra-regional specializations of neuronal clusters for many brain structures [16, 17, 18]. Using a simple average functional signal from an anatomically defined brain region over simplifies this complexity and risks missing information related to inter-regional variations in function.

To learn the location and number of functional nuclei in areas constrained by an anatomical atlas, we observe that if an anatomical region has many nuclei, we would expect to find multi-modal signals within the functional imaging data. Further, because functional imaging data is always smoothed with a Gaussian smoothing kernel during image formation, we would expect any signal from the nuclei, as measured by functional imaging, to be smoothly varying in space with a functional peak spatially close to the center-of-mass of the nuclei and an area roughly consistent with the extent of the cells composing the nuclei. Given these observations, we propose to model the functional clusters using finite Gaussian mixture models [19]. Based on data from neuroanatomic studies, there may be multiple nuclei within many anatomically-defined regions (e.g. Toncray and Krieg [20] discuss twenty-six nuclei of the human thalamus) but we do not know how many of these nuclei result in functional signals detectable in the acquired data. We can use this information to limit the number of clusters we search for and use model selection methods to identify how many clusters are needed to reasonably explain our observed functional signals.

We begin by thinking about the imaging data as being composed of a set of voxels, each being identified by its three-dimensional coordinate  $x_i = (p_1, p_2, p_3)$  within the image. For a given voxel coordinate,  $x_i$ , the acquired functional imaging data, is quantified to counts per unit time and provides a non-negative, real-valued proxy  $I(x_i) \in \mathbb{R}_{\geq 0}$  corresponding to the number of detected events by the imaging system at that voxel coordinate. We make the simplifying assumption that each detected event is independent from all other detections and identically distributed. This assumption implies that the functional imaging signal  $I(x_i)$  at each voxel (i.e., the number of detected events) can be modeled as  $I(x_i)$  independent observations at voxel  $x_i$ . For each region  $r$  defined in an anatomical atlas, the set of  $N^r$  voxel coordinates contained within the boundaries of the atlas region  $r = \{x_1, \dots, x_{N(r)}\}$  specify the voxels assigned to the anatomical region.

Each voxel is a member of only one anatomical region. The log-likelihood of all the detected events within a region is then modeled by a K-component Gaussian mixture model:

$$\ln p(X^r | \Theta^{(r)}) = \sum_{i=1}^{N^{(r)}} I(x_i) \ln \left\{ \sum_{k=1}^K w_k^{(r)} \mathcal{N}(x_i | \mu_k^{(r)}, \Sigma_k^{(r)}) \right\} \quad (\text{A.1})$$

where  $\Theta^{(r)} = \{w_1^{(r)}, \dots, w_K^{(r)}, \mu_1^{(r)}, \Sigma_1^{(r)}, \dots, \mu_K^{(r)}, \Sigma_K^{(r)}\}$  are the model parameters,  $X^r$  is the set of all voxel coordinates contained in region  $r$ , and  $I(x_i)$  is the number of detected events at the voxel coordinate  $x_i$  for all  $N^{(r)}$  coordinates. Parameters  $\mu_k^{(r)}$  and  $\Sigma_k^{(r)}$  represent the mean spatial location and covariance matrix of the  $k^{th}$  cluster in region  $r$ . The mixing weights  $w_k^{(r)}$  represent the probability that  $x_i$  was generated by component  $k$  and are subject to the constraints  $\sum_{k=1}^K w_k^{(r)} = 1$  and  $0 \leq w_k^{(r)} \leq 1$  to be valid probabilities. The unknown parameters  $\Theta^{(r)}$  can be learned using the iterative expectation maximization (EM) algorithm and maximized in closed form [21, 22]. Although the algorithm is guaranteed to converge, it is not guaranteed to converge to a global maximum of the likelihood function and is initialization dependent. To initialize the mixture model we use the kmeans++ algorithm of Arthur and Vassilvitskii [23], an extension to the K-means algorithm. K-means is a clustering technique which finds a partitioning of the data into  $K$  clusters that minimizes the sum of the squared distances between each data point and its closest cluster center and is routinely used to initialize Gaussian mixture models. The kmeans++ algorithm is a simple extension that uses a random seeding technique which has been shown to improve speed and accuracy with respect to the recovering the true cluster centers. To initialize our Gaussian mixture models using this technique, we choose an initial cluster center from the data for a particular anatomical region with the highest number of detected events. We then choose the next cluster center at random from the remaining data points with probability proportional to the distance of each data point to the closest previously selected cluster center, weighted by the number of detected events. This step gives higher probability to data points farther away from the previously selected cluster centers, weighted by the number of detections at those points. The procedure continues until all initial  $K$  cluster centers have been selected. We then run the k-means algorithm to convergence and use the assignments of the voxels to the  $K$  clusters to initialize our Gaussian mixture model means and covariances.

The final step in the clustering model is to select an appropriate setting for the number of clusters  $K$  for each region. We cannot simply choose the model that explains the data best because more complex models (i.e., larger  $K$ ) will always explain the data better than a simpler model with fewer clusters. Among the various approaches to model selection, penalized likelihood methods have been shown to be competitive, simple to implement in practice, and to work well when applied to mixture models [24]. Here we choose the Bayesian information criterion (BIC) developed by Schwarz et al. [25] for model selection, which has been shown to be a consistent estimator for independent and identically distributed observations in linear exponential family models such as the Gaussian mixture model [26]. By choosing the model with the minimum BIC score, one is attempting to select the candidate model with the highest Bayesian posterior probability. The BIC score

penalizes the model likelihood and tends to favor simple models over more complex ones. The BIC score for our mixture model within a brain region is given by:

$$BIC\left(\Theta^{(r)}\right)=-2 \ln p\left(X \mid \Theta^{(r)}\right)+\rho^{(r)} \ln \left(\sum_{i=1}^{N^{(r)}} I\left(x_i\right)\right) \quad (\text{A.2})$$

where  $\rho^{(r)}$  is the number of parameters in the model for region  $r$  and is used to penalize the likelihood relative to model complexity. In our model we are learning  $K$  three-dimensional Gaussian distributions for each region. For each Gaussian distribution of dimension  $D=3$ , we need  $D/2(D+1)$  parameters to specify the symmetric covariance matrix and  $D$  parameters to specify the mean. There are  $K$  mixing weights for the  $K$ -component Gaussian mixture model, resulting in  $K * D/2(D+3) + K$  parameters. After learning the mixture models for each setting of  $K$ , we compute the BIC score, penalizing the likelihood scores by the number of parameters estimated. We select the best model as the one with the lowest BIC score, which implies the model with the highest relative data likelihood after penalizing for model complexity. Once the best models have been selected for each brain region, we compute values for each cluster by averaging over the number of detections at each voxel, weighted by their cluster membership probabilities. We use these cluster averages, computed for each subject, in the sparse inverse covariance models.

## Appendix B. Inverse covariance model results without clustering

The results showing the average area under the TP/FP edge curves (AUC) across the subsets, by sample sizes, using the gold standard for parameter settings as described in section 3.0.1 but without the clustering model are shown in Table B.5. The functional data was averaged in each region of interest instead of using the results from the Gaussian mixture model described in Appendix A. The AUCs are statistically lower in the GL model (i.e. negative t-scores) than both the FGL and GGL models at data set sizes from 50-2000, with the moderate sample sizes showing the largest decreases. These results are consistent with those when using the clustering model, yet the differences between inverse covariance models are lower in magnitude and the AUCs are slightly higher, likely due to the smaller number of variables relative to sample size.

The results showing the average minimum SSEs across the subsets, by sample sizes, using the gold standard for parameter settings as described in section 3.0.1 but without the clustering model are shown in Table B.6. Again, the results are consistent with those presented using the clustering model, but the SSEs are less accurate over the range of sample sizes likely related to averaging functional signals with noise in the larger regions of interest.

Table B.5: TP/FP Edge Area Under Curve (AUC) T-test comparison by model and sample size using gold standard for parameter settings and simple averaging of functional signal in AAL regions of interest

	GL	FGL	GGL	GL-FGL	GL-GGL	FGL-GGL
N	$\overline{AUC} \pm std$	$\overline{AUC} \pm std$	$\overline{AUC} \pm std$	T(P Bonferroni)	T(P Bonferroni)	T(P Bonferroni)
2500	0.840±0.016	0.851±0.016	0.856±0.011	-1.538(1.00E+00)	-2.603(7.01E-01)	0.814(1.00E+00)
2000	0.823±0.009	0.842±0.011	0.842±0.010	-4.157( <b>2.31E-02</b> )	-4.388( <b>1.38E-02</b> )	-0.043(1.00E+00)
1500	0.806±0.011	0.828±0.010	0.824±0.007	-4.517( <b>1.04E-02</b> )	-4.421( <b>1.29E-02</b> )	-0.773(1.00E+00)
1000	0.772±0.009	0.801±0.011	0.796±0.009	-6.398( <b>1.96E-04</b> )	-5.937( <b>5.00E-04</b> )	-1.184(1.00E+00)
750	0.751±0.009	0.782±0.009	0.775±0.007	-7.511( <b>2.32E-05</b> )	-6.623( <b>1.26E-04</b> )	-1.933( <b>1.00E+00</b> )
500	0.720±0.011	0.753±0.012	0.742±0.009	-6.426( <b>1.86E-04</b> )	-5.053( <b>3.23E-03</b> )	-2.192(1.00E+00)
250	0.656±0.011	0.682±0.014	0.675±0.013	-4.600( <b>8.67E-03</b> )	-3.505( <b>9.85E-02</b> )	-1.195(1.00E+00)
100	0.617±0.011	0.638±0.011	0.631±0.009	-4.352( <b>1.50E-02</b> )	-3.177(2.04E-01)	-1.531(1.00E+00)
50	0.576±0.014	0.604±0.013	0.597±0.011	-4.502( <b>1.07E-02</b> )	-3.708( <b>6.27E-02</b> )	-1.189(1.00E+00)
25	0.560±0.014	0.576±0.011	0.574±0.008	-2.869(3.98E-01)	-2.622(6.75E-01)	-0.591(1.00E+00)
10	0.547±0.013	0.559±0.015	0.552±0.010	-1.878(1.00E+00)	-0.945(1.00E+00)	-1.199(1.00E+00)

Table B.6: Minimum average sum of squared errors (SSE) by model and sample size using gold standard for parameter settings and simple averaging of functional signal in AAL regions of interest

	GL	FGL	GGL	GL-FGL	GL-GGL	FGL-GGL
N	$\overline{SSE} \pm std$	$\overline{SSE} \pm std$	$\overline{SSE} \pm std$	T(P Bonferroni)	T(P Bonferroni)	T(P Bonferroni)
2500	5.94±5.08	7.13±6.11	9.59±8.30	-0.472(1.00E+00)	-1.184(1.00E+00)	-0.753(1.00E+00)
2000	7.69±8.88	9.07±7.79	14.36±12.14	0.397(1.00E+00)	-0.800(1.00E+00)	-1.159(1.00E+00)
1000	17.39±14.67	11.03±9.49	22.79±19.20	1.150(1.00E+00)	-0.707(1.00E+00)	-1.736(1.00E+00)
750	21.90±18.63	13.19±11.34	25.88±22.59	1.263(1.00E+00)	-0.430(1.00E+00)	-1.588(1.00E+00)
500	33.68±28.56	19.42±17.05	23.84±20.19	1.356(1.00E+00)	0.890(1.00E+00)	-0.529(1.00E+00)
250	44.23±37.35	42.54±37.06	41.11±34.53	0.102(1.00E+00)	0.194(1.00E+00)	0.089(1.00E+00)
100	78.07±67.43	60.83±51.49	69.76±60.06	0.643(1.00E+00)	0.291(1.00E+00)	-0.357(1.00E+00)
50	116.36±98.85	94.76±79.63	101.89±85.90	0.538(1.00E+00)	0.350(1.00E+00)	-0.192(1.00E+00)
25	188.52±159.39	165.43±140.50	180.41±153.59	0.344(1.00E+00)	0.116(1.00E+00)	-0.228(1.00E+00)
10	241.64±204.92	224.41±189.45	231.57±195.92	0.195(1.00E+00)	0.112(1.00E+00)	-0.083(1.00E+00)